arXiv:1612.00534v1 [cs.CV] 2 Dec 2016

Object Detection via End-to-End Integration of Aspect Ratio and Context Aware Part-based Models and Fully Convolutional Networks

Bo Li[1], Tianfu Wu[3,*], Shuai Shao[1,2,†], Lun Zhang[1] and Rufeng Chu[1]
[1]YunOS BU, Alibaba Group [2]Jilin University
[3]Department of Electrical and Computer Engineering and the Visual Narrative Cluster, North Carolina State University
{*shize.lb,lun.zhangl, rufeng.churf*}*@alibaba-inc.com, tianfu_wu@ncsu.edu, shaoshuai2113@mails.jlu.edu.cn*

Abstract

This paper presents a framework of integrating a mixture of part-based models and region-based convolutional networks for accurate and efficient object detection. Each mixture component consists of a small number of parts accounting for both object aspect ratio and contextual information explicitly. The mixture is category-agnostic for the simplicity of scaling up in applications. Both object aspect ratio and context have been extensively studied in traditional object detection systems such as the mixture of deformable part-based models [13]. They are, however, largely ignored in deep neural network based detection systems[17, 16, 39, 8]. The proposed method addresses this issue in two-fold: (i) It remedies the wrapping artifact due to the generic RoI (region-of-interest) pooling (e.g., a 3×3 *grid) by taking into account object aspect ratios. (ii) It models both global (from the whole image) and local (from the surrounding of a bounding box) context for improving performance. The integrated framework is fully convolutional and enjoys end-to-end training, which we call the aspect ratio and context aware fully convolutional network (ARC-FCN). In experiments, ARC-FCN shows very competitive results on the PASCAL VOC datasets, especially, it outperforms both Faster R-CNN [39] and R-FCN [8] with significantly better mean average precision (mAP) using larger value for the intersection-over-union (IoU) threshold (i.e.,* 0.7 *in the experiments). ARC-FCN is still sufficiently efficient with a test-time speed of* 380*ms per image, faster than the Faster R-CNN but slower than the R-FCN.*

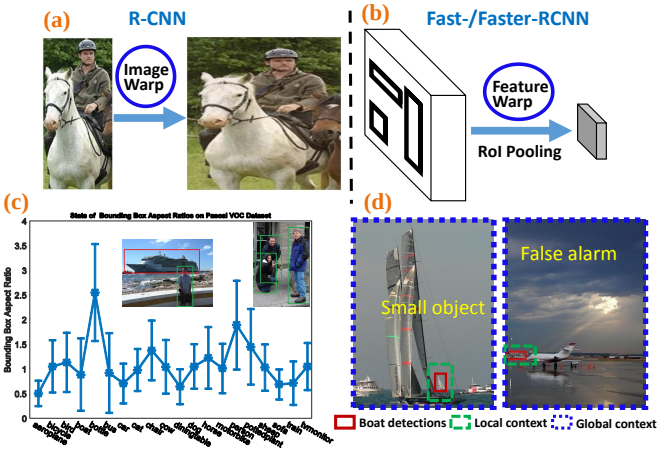

Figure 1. Motivation of the proposed method. It remedies the wrapping artifacts in state-of-the-art R-CNN object detection systems (see (a) and (b)) by integrating aspect ratio in the RoI pooling to respect the underlying distribution of aspect ratios (see plots and examples in (c)). It also takes into account both global and local contextual information (see the example of boat detection in (d)). See text for details.

1. Introduction

1.1. Motivation and Objective

We have witnessed a critical shift in the literature of object detection from explicit models such as the mixture of deformable part-based models (DPMs) [13] and its many variants, and hierarchical and compositional And-Or graph models [27], to less transparent but much more accurate deep neural network based approaches [26, 17, 16, 39, 21, 46, 38, 31, 8]. A popular framework of deep neural networks in object detection is the region-based convolutional neural networks (R-CNN) [17], which consists of two components: (i) A region proposal component is used to reduce the number of candidates to be classified since the sliding window technique is practically prohibitive. The proposals

*Tianfu Wu is the corresponding author.
†This work was done when Shuai Shao was an Intern at YunOS BU, Alibaba Group.

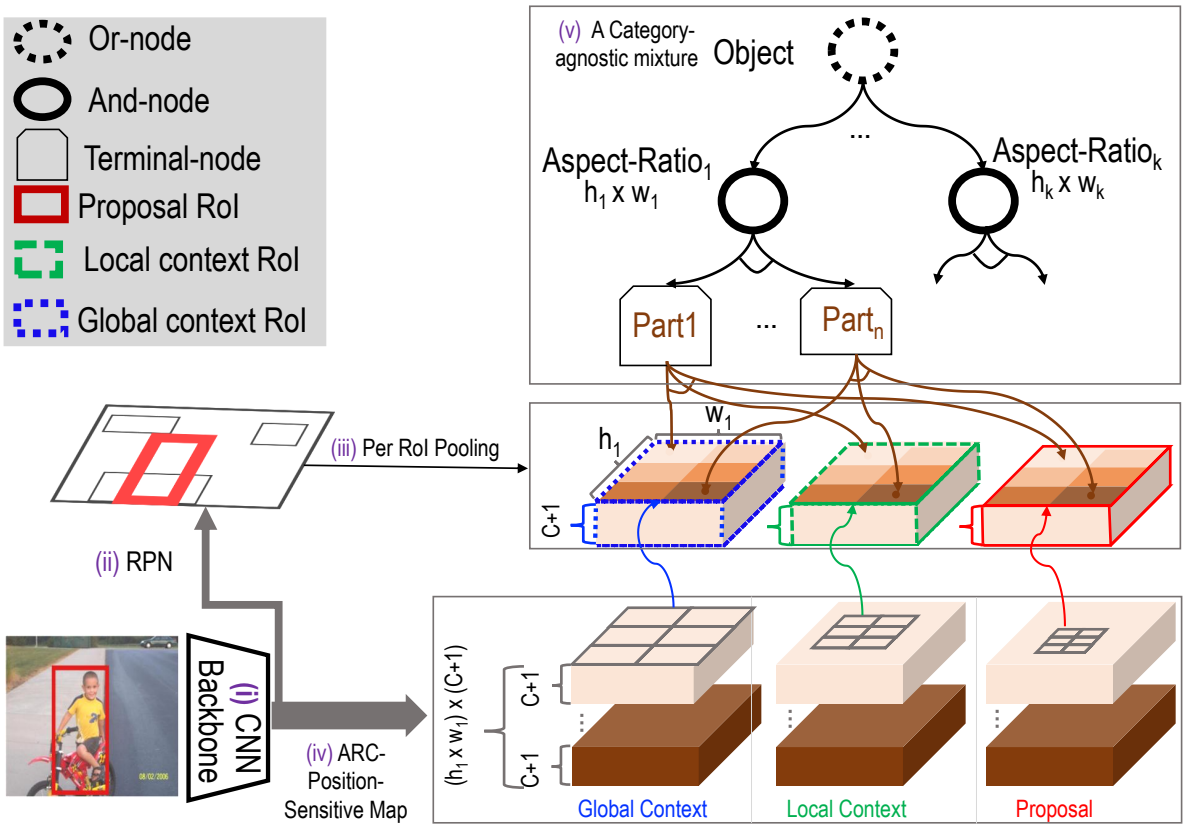

Figure 2. Illustration of the proposed method. It integrates a category-agnostic mixture of aspect ratio and context aware part-based models (represented by a simple And-Or tree as illustrated at the right-top) with state-of-the-art region-based fully convolutional networks [8] (see an illustration at the bottom). For clarity, we only show the work flow for one aspect ratio and the others are similar. See text for details. (Best viewed in color.)

are generated either by utilizing off-the-shelf objectness detectors such as the selective search [49], Edge Boxes [57] and BING [7] or by learning an integrated region proposal network (RPN) [39] in an end-to-end way. (ii) A prediction component is used to classify all the proposals and regress bounding boxes of the classified object candidates. To accommodate different shapes of the proposals, either raw image patches are wrapped to the same canonical size as done in the R-CNN [17] (see Fig. 1 (a)) or later on a more effective feature wrapping operator is designed based on the region-of-interest (RoI) pooling [16] (see Fig. 1 (b)) to compute equally-sized feature maps. These wrapping artifacts are purely caused by design choices for the simplicity of applying deep neural networks and for the practical consideration of affordable training and testing time complexities. They are less elegant than traditional explicit models in term of respecting the underlying distribution of object shapes (see Fig. 1 (c)), although they obtain much better performance.

This paper is motivated by two straightforward and intuitive questions: What are there that have shown performance improvement in the traditional object detection systems, but are largely ignored in the R-CNN framework and its variants? And, would they also improve the detection performance of state-of-the-art R-CNN object detection systems if they were integrated properly? Among many others, this paper addresses two aspects: one is to remedy the feature wrapping artifacts by accounting for aspect ratios in RoI pooling explicitly and the other is to take into account multi-scale contextual information, while most of the deep neural network based detection systems take advantage of global context only. Fig. 1 (d) illustrates the different roles played by local and global context. We can see the whole image of the airport scenario is quite similar to the one of the sea. In this case, global context will be confused, but the local context will be still helpful as it shows surrounding the object is a building, which most likely not a detection of a boat. To that end, this paper presents an end-to-end integration of a category-agnostic mixture of aspect ratio and context aware part-based models and state-of-the-art region-based fully convolutional networks, and it shows the integration of aspect ratio and context can benefit the widely used R-CNN framework significantly. We call our model the ARC-FCN (Aspect Ratio and Context aware

Fully Convolutional Network) and give a brief overview of ARC-FCN in the following section.

1.2. Method Overview

Fig. 2 shows overall architecture of the proposed ARC-FCN. Object detection via ARC-FCN in an image consists of two stages: generating category-agnostic bounding box proposals (or so-called objectness detection) and classifying each proposal into one of the $C+1$ categories (e.g., the $C=20$-class in the PASCAL VOC dataset and a background class) in terms of prediction scores with the bounding box regressed. The former is implemented by the RPN component and the latter is done through the RoI-wise prediction based on the category-agnostic mixture model.

The ARC-FCN consists of five components: (i) a fully convolutional network backbone using state-of-the-art network architecture such as the residual net [21] (with the average pooling layer and fully connected layer removed). (ii) a region proposal network (RPN) which computes bounding box proposals for all the C categories and thus it is category-agnostic. (iii) a ARC-position-sensitive map which is built on top of the position-sensitive map proposed in the R-FCN [8]. (iv) a per RoI pooling layer which extracts features for each cell in a RoI from the ARC-position-sensitive map. (v) a category-agnostic mixture of ARC-aware part-based models which computes the prediction score and a set of regression parameters for each RoI from the RPN.

The RoI pooling plays the central role in the architecture in similar spirit to the Faster R-CNN [39] and the R-FCN [8]. To address the issue of feature wrapping artifacts caused by using a single tiling rule in RoI pooling, we propose to use a mixture of RoI tiling rules each of which accounts for a different aspect ratio, in similar spirit to the mixture of DPMs [13] and the Or-nodes in an And-Or graph [55].

Unlike R-FCN [8] which predicts the class label for each proposal based on voting after RoI pooling, the parts in our mixture model have their own parameters. The ARC-position-sensitive map extends the position-sensitive map proposed in the R-FCN [8] in a straightforward way. The feature map for a part is pooled from the ARC-position-sensitive map using the same RoI pooling scheme as in [8, 16, 39]. Basically, the category-agnostic mixture of ARC-aware part-based models is the simplest part-based models built on top the position-sensitive feature maps.

Our ARC-FCN is fully convolutional and can be trained end-to-end following the similar procedure in the R-FCN [8] and the Fast/Faster R-CNN [16, 39]. We note that we use class-agnostic aspect ratios for simplicity and strictly speaking we do not eliminate the feature wrapping artifacts. The proposed method is, however, generic and can be easily extended to handle class-specific aspect ratios at the cost of model complexity and training and testing time.

In experiments, we test our model on the Pascal VOC 2007 and 2012 datasets [12]. Using the 50-layer and 101-layer Residual Net (ResNet-50 and ResNet-101) [21] as the backbone, our ARC-FCN yields consistently better performance in terms of average precision than the Faster R-CNN [39] and R-FCN [8] counterparts on both overlap requirements of ≥ 0.5 and ≥ 0.7 under the Pascal VOC evaluation protocol. Meanwhile, the test-time speed is about 380ms per image, which is faster than the vanilla Faster R-CNN. These results show the effectiveness and efficiency of our ARC-FCN.

2. Related Work and Our Contributions

CNN-based Detection Models. Object detection has been greatly improved by convolutional neural networks (CNN) on both accuracy and speed [11, 47, 17, 42, 17, 16, 39, 8, 38, 31]. Currently, there are two most influential research streams: a) along the first stream is the R-CNN style models [17, 25, 16, 39, 8], which is region-based, consisting of a region proposal module and a recognition module. Those works all rely on pre-computed object proposals to predict the position of each underlying object; b) along the second stream is the YOLO-style models [38, 31], which removes the proposal generating module and predicts object positions directly. Though YOLO-style models are faster than R-CNN style models, the R-CNN style models seems more accurate, as they still perform leading accuracies on popular benchmarks [12, 40, 30].

R-CNN style models often use either image warp [17, 54, 56], or feature warp [25, 16, 39] to transform the whole image feature to per RoI feature. Usually, the transformed per RoI feature are equally-sized. Though simplicity, the choice of those warping artifacts is a compromised consideration of effectiveness and affordable computational costs. As studied in [51], this is less elegant than traditional models, e.g., DPM [13, 18, 1, 2], AOG [55, 27].

Mixture Modelling. Many traditional works [13, 18, 1, 2, 53, 34, 52, 32, 37, 27] explicitly model the intra-class or sub-category variations by mixture modelling. Specifically, [13, 18] model mixtures by the aspect ratio of object bounding box. This is a simple yet effective way to quantify the visual space. [32] models mixtures by 2D viewpoints, [36, 5] model mixtures by 3D viewpoints, [37, 35, 27] model mixtures by occlusion extent of an object. In [1, 34, 52], mixtures are discovered by feature clustering. In [53, 6], mixtures are modelled based on the relative distance of neighbouring joints.

In the literature, [19, 51] studied integrating mixture models with CNN framework, but the performance is still worse than region-based CNN object detectors.

Contextual Modelling. A lot of traditional detection models [10, 27, 20, 23, 48] exploit the role of context. Re-

cently, [50] utilized the global context and object-to-object pairwise relationships for person head detection. [29] utilized hierarchical context for human attribute recognition, [28] used multi-level context for person recognition, [43] used contextual priming to improve Faster R-CNN models. For ResNets-based Faster R-CNN model [21], only global context was utilized, and it introduce expensive computational costs.

Works on Accurate Localization. There are many works related to accurate object localization. [14] incorporated segmentation cues with DPM, [9] utilized color and edge features to refine the bounding box coordinates of DPM detections, [41] used the height prior of an object. [54] used Bayesian optimization to refine the bounding box proposals and trained the CNNs with structured loss on inaccurate localization. [15] assigned probabilities on each row and column of a search region to get accurate object positions.

We believe explicitly modelling mixtures of aspect ratio and contextual information as well as many other factors are universally useful in object detection, especially when integrated with powerful deep neural networks end-to-end. And this paper makes the following three contributions to object detection.

- It proposes a simple yet effective framework to integrate explicit models (e.g., category-agnostic aspect ratio and context used in this paper) with fully convolutional networks for accurate and efficient object detection.
- It proposes the aspect ratio, context, and position sensitive RoI pooling layers to bridge explicit models and deep neural networks.
- It obtains state-of-the-art detection performance on the PASCAL VOC 2007 and 2012 datasets w.r.t. comparable baselines.

3. The Proposed Model

In this section, we present details of the proposed ARC-FCN (see Fig. 2).

Denote by Λ an image lattice and by I_Λ an image defined on the lattice Λ. We formulate the two stages in object detection via ARC-FCN in the following.

The RPN subnetwork, denoted by $r(I_\Lambda; \Theta_{RPN})$, computes a set of category-agnostic bounding box proposals (foreground vs background),

$$r(I_\Lambda; \Theta_{RPN}) = \{(B_i, t_i, l_i, p_i); i = 1, \cdots, M\} \quad (1)$$

where Θ_{RPN} is the parameters including parameters in the convolution network backbone, $l_i \in \{0, 1\}$, $t_i = (t_i^x, t_i^y, t_i^{wd}, t_i^{ht})$ a 4-d vector of bounding box regression parameters used to refine B_i and p_i (the so-called objectness probability). The total number M is determined by the size of the feature map of the last layer in the convolutional network backbone (see Fig. 2) and the number of translation-invariant anchor boxes (details are referred to the Faster R-CNN [39]). For a pair of (B, t), let $B = (x, y, wd, ht) \subseteq \Lambda$ where (x, y) represents the pixel coordinates of the bounding box center and (wd, ht) the width and height respectively. The refined bounding box B' is computed following the parameterization in Fast R-CNN [16]: $B' = (x', y', wd', ht')$ where $x' = t^x \cdot wd + x$, $y' = t^y \cdot ht + y$, $wd' = wd \cdot \exp(t^{wd})$ and $ht' = ht \cdot \exp(t^{ht})$.

Before feeding into the prediction subnetwork, the set of foreground proposals (i.e., $l_i = 1$) is post-processed with non-maximum suppression (NMS) using a predefined intersection-over-union (IoU) threshold τ_{RPN}. For notional simplicity, we write $B = (x, y, wd, ht)$ as a foreground bounding box proposal (i.e., a RoI to the prediction subnetwork).

The category-agnostic mixture model consists of an Object Or-node, a small number K of aspect ratio And-nodes (h_k, w_k) $(k = 1, \cdots, K)$ and $h_k \times w_k$ part Terminal-nodes for a corresponding And-node. The selection of aspect ratios and the number K are category-agnostic and takes into account their statistics in training data of all object categories (see also the ablation studies in the experiments). We will omit k when referring aspect ratios since they are functionally same. For a RoI B, the outputs of all nodes are in the same format consisting of two components: one is a vector of prediction scores $\{s^l(v); l \in \{0, 1, \cdots, C\}\}$ and the other is a 4-d vector of bounding box regression parameters $(t^x(v), t^y(v), t^{wd}(v), t^{ht}(v))$, where v represents a node in the mixture model.

An And-node A represents a SUM operator and we compute its prediction score $s^l(A) = \sum_{v \in ch(A)} s_i^l(v)$ where $ch(A)$ is the set of child nodes, and its bounding box regression $t^{reg}(A) = \sum_{v \in ch(A)} t^{reg}(v)$ where $t^{reg}() \in \{t^x(), t^y(), t^{wd}(), t^{ht}()\}$. An Or-node represents a MAX operator and the outputs of the object Or-node O is defined by: $s^l(O) = s_l(A^*)$ and $t^{reg}(O) = t^{reg}(A^*)$ where $A^* = \arg\max_A \max_l(s_l(A))$. Then we compute the prediction probability by normalizing $s_l(O)$ using the Softmax operator: $p^l(O) = \frac{s^l(O)}{\sum_{l'} s^{l'}(O)}$.

The class label of a RoI B is predicted based on $\ell = \arg\max_l p^l(O)$ and the refined bounding box B' is computed based on $t^{reg}(O)$ in the same way as the RPN. So, we have the detection results of an image I_Λ: $y = \{(B_j', t_j, \ell_j, p_j^{\ell_j}); j = 1, \cdots, M'\}$. To generate the final detection results, another round of NMS with a predefined IoU threshold τ_{ARC} is applied for y.

So, we can write the ARC-FCN as a structured output

prediction function,

$$f(I_\Lambda;\Theta) = (r(I_\Lambda), y) \tag{2}$$

where Θ collects all parameters in the ARC-FCN.

The scoring function of part Terminal-node. The proposed mixture model is of the simplest form and only part Terminal-nodes have parameters. The scoring function of a part Terminal-node v is defined by a linear function,

$$Score(v) = F_v \cdot \Theta_v \tag{3}$$

where F_v represents the feature map computed by RoI pooling which is a 3-d map with the number of channels $3(C+1)$, the width wd_v and height ht_v. The wd_v and ht_v are the template width and height of a part Terminal-node. The number of channels is $3(C+1)$ since the features are pooled from the RoI itself, the local surrounding context and the global context (to be specified below). Θ_v are the parameters being a 4-d map containing a number $(C+1)+4$ of 3-d maps of the same dimensions as F_v. So the outputs of a part Terminal-node consist of $C+1$ classification scores and a 4-d vector for bounding box regression.

The ARC-position-sensitive score map. To extract the feature map F_v for a Terminal-node v, we compute the ARC-position-sensitive score maps for each aspect ratio And-node, which consist of three maps, $\mathbb{F}_{roi}$, $\mathbb{F}_{local}$ and $\mathbb{F}_{global}$, denoting the position-sensitive score map for the RoI itself, the local context and the global context respectively. The three maps have the same dimensions each of which is a 3-d map with the number of channels $h \times w \cdot (C+1)$, the width W and height H, where the number of channels is computed as such since we have $h \times w$ positions due to the tiling based on the aspect ratio and each position accounts for $(C+1)$ categories (see Fig. 2).

The RoI pooling. Given a RoI B, denote by b_{roi}, b_{local} and b_{global} the transformed bounding boxes of B in the three feature maps ($\mathbb{F}_{roi}$, $\mathbb{F}_{local}$ and $\mathbb{F}_{global}$) respectively, which map the RoI from the pixel coordinates to the coordinates in the feature maps. b_{local} is centered at the same position as b_{roi} with the side length increased by a factor λ ($\lambda = 1.5$ in our experiments). $b_{global} = (0, 0, W, H)$ is the same for all RoIs. Th same RoI pooling procedure is done for the three bounding boxes b_{roi}, b_{local} and b_{global}. Without loss of generality and notional confusions, consider the pooling in a transformed bounding box $b = (x, y, wd, ht)$ with left-top coordinates (x, y), width wd and height ht, we first tile b into $w \times h$ cells based on the aspect ratio. Each cell $b_{i,j} = (x + (j-1)\lfloor \frac{wd}{w} \rfloor, y + (i-1)\lfloor \frac{ht}{h} \rfloor, \lfloor \frac{wd}{w} \rfloor, \lfloor \frac{ht}{h} \rfloor)$ is linked by a part Terminal-node v for grounding ($1 \le i \le h, 1 \le j \le w$ and the height of cells at the bottom row and the width of cells at the right-most column will compensate the rounding effect accordingly). The feature map of a part Terminal-node F_v is computed using average pooling from the three ARC-position-sensitive maps, as done in the R-FCN [8], so we have $wd_v = ht_v = 1$ in Eqn. 3 and F_v becomes a $3(C+1)$-d vector. We have,

$$F_v = (F_v^{roi}, F_v^{local}, F_v^{global}), \tag{4}$$

and $F_v^{roi}(c) = \sum_{(\mathbf{x},\mathbf{y}) \in b_{i,j}} \mathbb{F}_{roi}(c', \mathbf{x}, \mathbf{y}) / |b_{i,j}|$ where $c' = ((i-1) \cdot w + (j-1)) \cdot (C+1) + c$ is the position-sensitiv mapping between the channel index of the ARC-position-sensitive feature map and that of the feature map of a part Terminal-node. In the same way, we can compute F_v^{local} and F_v^{global}.

4. Parameter Learning

In this section, we first pose the learning of ARC-FCN as a multi-task learning problem which can be trained end-to-end. Then we present an iterative training procedure for efficiency in practice.

Training Data. Denote the set of N training images with bounding box annotations by,

$$D = \{(I_i, y_i^* = \{B_{i,j}^*, \ell_{i,j}^*\}_{j=1}^{m_i}); i = 1, \cdots, N\}, \tag{5}$$

where each image I_i consists of m_i objects with bounding box $B_{i,j}^*$ and ground-truth class label $\ell_{i,j}^* \in [1, C]$.

Since the objective of the RPN subnetwork is to generate bounding box proposals to reduce the computational burden for the prediction subnetwork, to train the RPN subnetwork, we augment the training dataset D with category-agnostic annotations following the method proposed in the Faster R-CNN [39]. The basic idea is to use sparse sample of all possible sliding windows guided by pre-specified anchor boxes which enable translation-invariant parameterization. For an image $I_i \in D$, we have the augmented annotations,

$$\Omega_i = (B_{i,k}^*, t_{i,k}^*, l_{i,k}^*), \tag{6}$$

where $k = m_i + 1, \cdots, m_i + M_i$ and $l_{i,k}^* \in \{0, 1\}$, $l_{i,k}^* = 0$ means that $B_{i,k}^*$ is treated as background and $l_{i,k}^* = 1$ means that $B_{i,k}^*$ is treated as foreground proposal, $t_{i,k}^*$ is the bounding box regression parameters computed for foreground $B_{i,k}^*$ w.r.t. the $B_{i,j}^*$'s. Details of generating the augmented annotations are referred to [39].

Loss Function. The loss function of ARC-FCN consists of two components: a loss function for the RPN subnetwork based on the augmented annotation Ω_i and the RPN prediction $r(I_i)$ in Eqn. 1, and a loss function for the prediction subnetwork based on the annotation y_i^* and the detection output y_i in Eqn. 2. For a training sample $(I_i, y_i^*, \Omega_i) \in D$, we define the loss by,

$$\begin{aligned} L_{\text{ARC-FCN}}(f(I_i;\Theta), y_i^*, \Omega_i) =& L_{RPN}(r(I_i), \Omega_i) + \\ & L_{Pred}(y_i, y_i^*) \end{aligned} \tag{7}$$

Both $L_{RPN}(r(I_i), \Omega_i)$ and $L_{Pred}(y_i, y_i^*)$ consist of a classification loss term and a regression loss term.

Loss function for the RPN subnetwork $L_{RPN}(r(I), \Omega)$ is defined by,

$$L_{RPN}(r(I_i), \Omega_i) = \frac{1}{M} \sum_k L_{cls}(p_{i,k}, l_{i,k}^*) + \lambda \cdot \frac{1}{M} \sum_k l_{i,k}^* \cdot L_{reg}(t_{i,k}, t_{i,k}^*) \quad (8)$$

where the regression loss is computed for foreground bounding boxes only and $L_{reg}(\cdot, \cdot)$ is defined by certain robust loss such as the smooth L_1 loss proposed in the Fast R-CNN [16]. In practice, since Ω_i is relatively large, the number of positives and the number of negatives used in a mini-batch are predefined.

Loss function for the prediction subnetwork $L_{Pred}(y_i, y_i^*)$ is defined similarly,

$$L_{Pred}(y_i, y_i^*) = \frac{1}{M'} \sum_j L_{cls}(p_{i,j}^{\ell_{i,j}}, l_{i,j}^*) + \lambda \cdot \frac{1}{M'} \sum_j l_{i,j}^* \cdot L_{reg}(B_{i,j}, B_{i,j}^*) \quad (9)$$

where $l_{i,j} \in [0, C]$, $B_{i,j}$ denote the predicted class label and bounding box for the j-th object on image I_i, respectively.

Cascade and iterative training. Given the training dataset D, we first train an RPN to get the initial set of proposals P^0. Then we train the prediction subnetwork with P^0 fixed. We iterate this procedure: during each iteration t, the prediction subnetwork is updated by the learning process based on image set D and proposals P^{t-1}, then proposals will be updated by model testing. In practice, we find the model is converged after 2 iterations (excluding the first RPN training).

Implementation details. ARC-FCN is implemented using the open source Caffe CNN library [24]. OHEM [44] is utilized for efficient training, 128 RoIs are selected for backpropagation [8]. To cope with small objects, the *À trous* algorithm [4, 33] is utilized to enlarge the last convolutional feature maps in ResNets. Training images are resized such that the min scale (shorter side of image) is 600 pixels, and the max scale (longer side of image) is 1000 pixels as [16]. The mini-batch size is 2. For each training step, we use a weight decay of 0.0005 and a momentum of 0.9. ARC-FCN is finetuned using a learning rate of 0.001 for the first $80k$ mini-batches and 0.0001 for another $30k$. For the first training step, we adopt 2 step training strategy without feature sharing between RPN and ARC-FCN, and this saved about $1/3$ time without losing performance.

PASCAL VOC 2007 (IoU ≥ 0.5)			
Method	training data	test times	mAP
R-FCN-Res50 [8]	07+12 *trainval*	0.12	77.4
R-FCN-Res50-ReIm	07+12 *trainval*	0.14	77.3
ARC-FCN-Res50	07+12 *trainval*	0.23	**80.2**
Faster R-CNN-Res101 [39]	07+12 *trainval*	0.42	76.4
R-FCN-Res101	07+12 *trainval*	0.17	79.5
R-FCN-Res101-ReIm	07+12 *trainval*	0.20	79.4
ARC-FCN-Res101	07+12 *trainval*	0.38	**82.0**

Table 1. mAP results with IoU ≥ 0.5 of Faster R-CNN, R-FCN and ARC-FCN on PASCAL VOC 2007 test set. ResNet-50 and ResNet-101 are used as the backbone architectures. Time is evaluated on a Nvidia K40 GPU.

5. Experiments

5.1. Pascal VOC Datasets

PASCAL VOC 2007 Testset. We first verify our method on the PASCAL VOC 2007 dataset [12]. This dataset has 20 object categories. For the training set, we use the union set of PASCAL VOC 2007 *trainval* (including $5,011$ images) and PASCAL VOC 2012 *trainval* (including $11,540$ images) following [39, 8]. As a common trick, these images will be flipped during training. For the testing set, we use the PASCAL VOC 2007 *test* (including $4,952$ images). During testing, non-maximum supression (NMS) is used to report the final results. We adopt the PASCAL VOC protocol [12], i.e., a detection is correct only if the intersection over union (IoU) of its bounding box and the groundtruth bounding box are equal or greater than 0.5, and evaluate our model by mean average precision (mAP). To verify our model on the ability of accurate localization, we also utilize the IoU ≥ 0.7 criterion. Our model is initialized by the ResNet-50 and ResNet-101 released with the paper of [21]. To cover most of aspect ratios among the PASCAL 20 object classes and get a balance of GPU memory and the computational efficiency, we use aspect ratios as $\{7 \times 7, 5 \times 10, 10 \times 5, 4 \times 12, 12 \times 4, 3 \times 12, 12 \times 3\}$ for ARC-FCN-Res50s, and $\{7 \times 7, 5 \times 10, 10 \times 5, 3 \times 12, 12 \times 3\}$ for ARC-FCN-Res101s. In this paper, we focus on investigating the role of bounding box aspect ratio and context. For simplicity, we don't use other training and testing tricks, e.g., multi-scale training, multi-scale testing, test with left-right flipped images [3]) and MS COCO data [30] for finetune.

Table 1 shows the IoU ≥ 0.5 results of our model and state-of-the-art Faster R-CNN [39] and R-FCN [8] models. For fairness, Faster R-CNN and R-FCN are also compared without using MS COCO data and multi-scale training. For the convenience of comparison with R-FCN on bounding box localization under the same setting, we also reimplement the R-FCN models as R-FCN-Res50-ReIm and R-FCN-Res101-ReIm. The mAP results is consistent with the original paper [8]. From Table 1, we can see our mod-

PASCAL VOC 2007 (IoU $\geq$ 0.7)							
Method	R-FCN-Res50-ReIm	R-FCN-Res101-ReIm	R-FCN-Res101 [8]	[54]	LocNet [15]	ARC-FCN-Res50	ARC-FCN-Res101
mAP	57.8	60.5	60.5	43.7	65.4	64.7	**68.2**

Table 2. mAP results with IoU $\geq$ 0.7 of ARC-FCN and state-of-the-art models on PASCAL VOC 2007 test set.

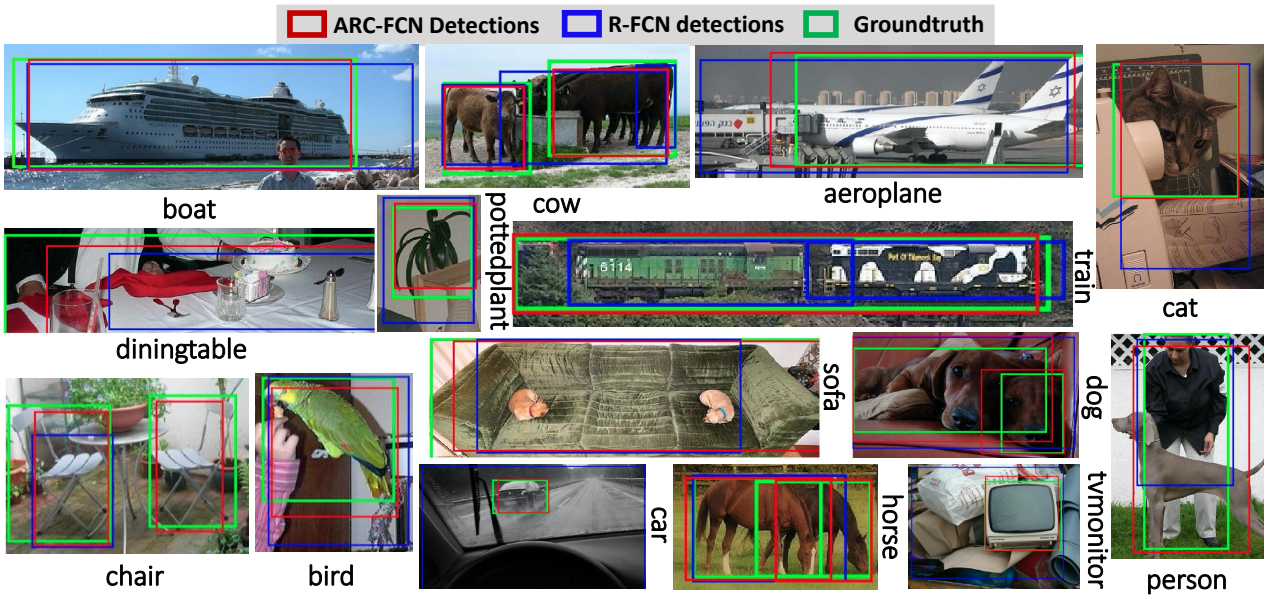

Figure 3. Sample detections of ARC-FCN-Res101 (red) and R-FCN-Res101 [8] (blue). For comparison, we also show the groundtruth bounding box (green). The score threshold is set to 0.6 for good visualization. Best viewed in color and zoom in.

els outperform R-FCN by 2.8 and 2.5 points in terms of mAP with ResNet-50 and ResNet-101, respectively, and 5.6 points over Faster R-CNN with ResNet-101, which verify the advantages of integrating aspect ratio and context.

Table 2 shows the IoU $\geq$ 0.7 results of our model and state-of-the-art models. Compared with Table 1, we can see ARC-FCN outperms R-FCN by a large margin 6.9 and 7.7 points in terms of mAP with ResNet-50 and ResNet-101, respectively. These results show that bounding box aspect ratio and context are more important on accurately localizing the positions of objects with diverse shapes, which is the motivation of this work. In Fig. 3, we show some qualitative results of ARC-FCN (red) and R-FCN (blue, using the ResNet-101 based model provided by the authors of [8]). For the convenience of comparison, we also show the annotations (green). From these results, we can see our model can localize objects more accurately.

From Table 2, we can see our model also outperforms previously state-of-the-art models that have better abilities on localizations than R-CNN style models [17, 16, 39]. In specific, ARC-FCN-Res50 is already comparative with the best LocNet in [15], and ARC-FCN-Res101 achieves 2.8 points improvement over the best LocNet in [15]. Since [54, 15] using different strategies on predicting precise object positions, i.e, applying Bayesian optimization and structural prediction [54], and modelling localization probabilities [15], they are complementary with our method.

Fig. 4 shows some qualitative results of our model on the PASCAL VOC 2007 testset. For visual convenience, different object classes are shown with different colors. On the left, we show some detection examples, we set a relatively high threshold of 0.45 for good displaying. From these results we can see, our model can fairly well localizing various objects. On the right, we show the typical failure examples. The failure cases are mainly due to occlusion, rough or partial detections and inter-class confusions (e.g., cat and dog, bus and train).

PASCAL VOC 2012 Testset. For PASCAL VOC 2012 benchmark [12], we use the union set of PASCAL VOC 2007 *trainval+test* (including 9,963 images) and PASCAL VOC 2012 *trainval* (including 11,540 images) as the training set following [39, 8]. These images are also flipped during training. The PASCAL VOC 2012 testset consists of 10,991 images. Training and testing strategies are the same as the previous experiment on PASCAL VOC 2007 testset.

Table 3 shows the IoU $\geq$ 0.5 results of our method and state-of-the-art Faster R-CNN [39, 21] and R-FCN [8]. For fair comparison, we reimplement R-FCN based on ResNet-101 without multi-scale training. We can see our method still outperforms them by 4.5 and 1.7 points of mAP respectively. We also list the testing time in Table 3. We can see ARC-FCN doesn't increasing the test time much, just

Figure 4. Qualitative Results of ARC-FCN-Res101 on Pascal VOC 2007. The score threshold is set to 0.45 for good visualization. Best viewed in color and zoom in.

PASCAL VOC 2012 (IoU ≥ 0.5)			
Method	training data	test times	mAP
Faster R-CNN-Res101	07+12 *trainval* + 07 *test*	0.42	73.8
R-FCN-Res101-ReIm† :	07+12 *trainval* + 07 *test*	0.17	76.7
ARC-FCN-Res101‡ :	07+12 *trainval* + 07 *test*	0.38	**78.4**

Table 3. mAP results with IoU $\geq$ 0.5 of Faster R-CNN, R-FCN and ARC-FCN on PASCAL VOC 2012 test set. Time is evaluated on a Nvidia K40 GPU. † : http://host.robots.ox.ac.uk:8080/anonymous/C3H0GM.html ‡ : http://host.robots.ox.ac.uk:8080/anonymous/WB5KF0.html

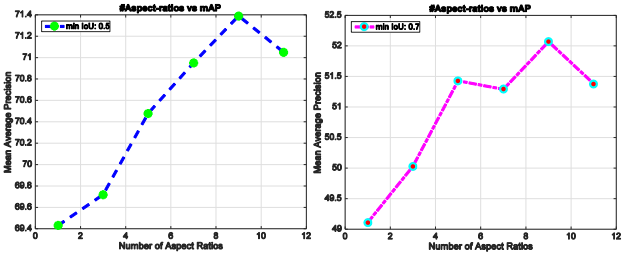

Figure 5. mAP results with different number of aspect ratio branches on PASCAL VOC 2007 test set.

about 380ms per image, though slower than R-FCN, still faster than Faster R-CNN.

5.2. Ablation Study

Effects of multi-aspect ratios. To investigate the role of aspect ratio modelling, we trained several ARC-FCNs but without context modelling, these models have different number of aspect ratios. We utilize the ResNet-50 for initialization and set the learning rate as 0.001 for the first 50k iterations training and 0.0001 for the next 20k iterations, with an effective mini-batch size of 1. For this pilot experiment, OHEM is not used. We investigate aspect ratios with five branches: a) $\{7\times7\}$, b) $\{7\times7, 5\times10, 10\times5\}$, c) $\{7\times7, 5\times10, 10\times5, 3\times12, 12\times3\}$, d) $\{7\times7, 5\times10, 10\times5, 4\times12, 12\times4, 3\times12, 12\times3\}$, e)$\{7\times7, 5\times10, 10\times5, 4\times12, 12\times4, 3\times12, 12\times3, 3\times15, 15\times3\}$. Fig. 5 shows the results, we can see increasing the number of aspect ratios will also increase the detection performance in general. However, when introducing too many asspect ratio branches, the performance may be saturate or dropped due to overfitting. This may be relieved by introducing more data or some effective regularization techniques (e.g., dropout [45]).

Effects of context. To investigate the impact of context, we trained several ARC-FCNs but without aspect ratio modelling. For this experiment, we utilize the ResNet-101 for weight initialization and set the learning rate as 0.001 for the first 80k iterations training, and 0.0001 for the next 30k iterations, with an effective mini-batch size of 2. OHEM is adopt for training. Table 4 shows the detailed results of ARC-FCN with/without local and global contexts, we can

Effects of Context without Aspect Ratio (PASCAL VOC 2007)		
Method	mAP (IoU ≥ 0.5)	mAP (IoU ≥ 0.7)
ARC-FCN-Res101-no-context	79.4	60.5
ARC-FCN-Res101-global-context	79.9	60.6
ARC-FCN-Res101-local-global-context[1]	80.4	63.3
ARC-FCN-Res101-local-global-context[2]	**80.7**	62.8

Table 4. mAP results with local and global context on PASCAL VOC 2007 test set. Local context is not used for bounding box regression in ARC-FCN-Res101-local-global-context[1], but used in ARC-FCN-Res101-local-global-context[2].

Iterative Two-step Training (PASCAL VOC 2007)		
Method	mAP (IoU ≥ 0.5)	mAP (IoU ≥ 0.7)
R-FCN-Res50-ReIm	77.3	57.8
ARC-FCN-Res50-Step1	**80.2**	60.2
ARC-FCN-Res50-Step2	80.1	**64.0**
R-FCN-Res101-ReIm	79.4	60.5
ARC-FCN-Res101-Step1	**82.0**	64.7
ARC-FCN-Res101-Step2	81.7	**68.2**

Table 5. Comparison of Two-Step ARC-FCN on mAP with both IoU ≥ 0.5 and IoU ≥ 0.7 on PASCAL VOC 2007 test set.

see both local context and global context can boost the detection performance. Local context is not used for bounding box regression in ARC-FCN-Res101-local-global-context[1], but used in ARC-FCN-Res101-local-global-context[2]. We can see local context can help localizing objects when min IoU is 0.5, but harms the result when the min IoU is 0.7. In our preliminary experiment (with ResNet-50), we find the global context harms the bounding box regression by 2.9 points, so we don't explore global context for bounding box regression here.

Effects of Two-Step ARC-FCN Training. Our ARC-FCN is actually an iterative framework, but 2-step training is enough on the Pascal VOC datasets. To analyze the effect of each training step, we compare the results of both IoU ≥ 0.5 and IoU ≥ 0.7 in Table 5. From the table, we can see the first step training already get powerful results on both IoU ≥ 0.5 and IoU ≥ 0.7. The second step training increase the results of IoU ≥ 0.7 much, but little on the results of IoU ≥ 0.5.

6. Conclusion

We presented the aspect ratio and context aware fully convolutional network (ARC-FCN) to integrate a mixture of part-based models and region-based convolutional networks for accurate and efficient object detection. The proposed method is fully convolutional and enjoys end-to-end training. It remedies the wrapping artifact due to the generic RoI (region-of-interest) pooling (e.g., a 3×3 grid) and models both global (from the whole image) and local (from the surrounding of a bounding box) context for improving performance. ARC-FCN outperforms both Faster R-CNN [39] and R-FCN [8] with significantly better average precision using larger value for IoU ≥ 0.7, meanwhile, it is still sufficiently efficient with a test-time speed of 380ms per image, faster than the faster R-CNN but slower than the R-FCN.

Acknowledgment: We thank Jifeng Dai et al for making the code of R-FCN [8] publicly available, and Chuan Yang, Yafei Song and Dong Yi for helpful discussions. T. Wu was supported by the ECE startup fund 201473-02119 at NCSU. B. Li was supported by the National Key Technology Research and Development Program of the Ministry of Science and Technology of China under Grant No. 2014BAK14B03.

A. Detailed Results on PASCAL VOC Datasets

In the article, we compared with Faster R-CNN [39] and R-FCN [8], with all methods using their pure forms (no multi-scale train/test and other tricks), to focus on the investigation of aspect ratio and context. We show the detailed results on PASCAL VOC Datasets [12] in the following.

A.1. Pascal VOC 2007 Test Set.

In this section, we show the detailed detection results of Faster R-CNN [39], R-FCN [8], and ARC-FCN on PASCAL VOC 2007 test set. For this dataset, the union set of PASCAL VOC 2007 *trainval* and PASCAL VOC 2012 *trainval* is used as the training set, and the PASCAL VOC 2007 *test* is used as the test set following [39, 8]. We adopt the PASCAL VOC evaluation protocol [12], and evaluate our model by mean average precision (mAP). Both intersection over union (IoU) ≥ 0.5 and ≥ 0.7 are utilized for PASCAL VOC 2007 (not apply for PASCAL VOC 2012), as detailed annotations are available only on PASCAL VOC 2007.

Table 6 shows the results with $IoU \geq 0.5$. Here, ARC-FCN is trained with One-Step training (excluding RPN training). We can see ARC-FCN outperforms the other 2 methods on all object categories. For some object classes, the improvement is quite impressive (above 3 points in terms of AP), e.g., areoplane, bike, bird, boat, chair, table, and person.

Table 7 shows the results with $IoU \geq 0.7$. Here, ARC-FCN is trained with Two-Step training (excluding RPN training). The model of R-FCN-Res101 is provided by the authors of [8]. As R-FCN is superior than Faster R-CNN, we just compare our method with R-FCN here. We can see ARC-RFCN outperforms R-FCN by a even larger margin across all object categories. For some object classes, e.g., aeroplane, bike, boat, bottle, chair, person, the improvements is above 10 points in terms of AP. These results verified the superior localization ability of ARC-FCN.

Besides, we also compare the detection diagnosis of R-FCN-Res101 and ARC-FCN-Res101 in Fig. 6. This is done

method	areo	bike	bird	boat	bottle	bus	car	cat	chair	cow	table	dog	horse	mbike	person	plant	sheep	sofa	train	tv	mAP
Faster-RCNN-Res101	79.8	80.7	76.2	68.3	55.9	85.1	85.3	89.8	56.7	87.8	69.4	88.3	88.9	80.9	78.4	41.7	78.6	79.8	85.3	72.0	76.4
R-FCN-Res101	82.5	83.7	80.3	69.0	69.2	87.5	88.4	88.4	65.4	87.3	72.1	87.9	88.3	81.3	79.8	54.1	79.6	78.8	87.1	79.5	79.5
ARC-FCN-Res101	**88.3**	**88.0**	**83.4**	**75.8**	**71.0**	**88.0**	**89.2**	**89.6**	**68.5**	**88.3**	**77.3**	**88.9**	**88.7**	**85.1**	**83.2**	**54.3**	**83.7**	**80.6**	**87.7**	**80.7**	**82.0**

Table 6. Detailed results with IoU ≥ 0.5 of Faster-RCNN, R-FCN and ARC-FCN on PASCAL VOC 2007 test set.

method	areo	bike	bird	boat	bottle	bus	car	cat	chair	cow	table	dog	horse	mbike	person	plant	sheep	sofa	train	tv	mAP
R-FCN-Res101	61.7	66.9	58.0	43.4	41.5	78.9	70.7	74.4	41.0	68.6	47.1	72.3	69.6	62.9	57.1	31.2	67.0	60.1	70.5	66.9	60.5
ARC-FCN-Res101	**78.1**	**76.3**	**66.6**	**54.5**	**58.2**	**80.4**	**78.8**	**78.2**	**51.0**	**70.9**	**55.5**	**74.1**	**75.9**	**69.0**	**67.1**	**39.5**	**74.8**	**64.2**	**76.7**	**74.2**	**68.2**

Table 7. Detailed results with IoU ≥ 0.7 of R-FCN and ARC-FCN on PASCAL VOC 2007 test set.

by applying the excellent detection analysis tool from [22]. For good visualization, we just show some typical object classes (i.e., aeroplane, bottle, boat) with large localization improvements by ARC-FCN. The first two rows show the result of R-FCN, and the last two rows show the result of ARC-FCN. From these plots, we can see the localization error of ARC-FCN is smaller than the one of R-FCN.

A.2. Pascal VOC 2012 Test Set.

In this section, we show the detailed detection results ($IoU \geq 0.5$) of Faster R-CNN [39], R-FCN [8], and ARC-FCN on PASCAL VOC 2012 dataset (see Table 8). For this dataset, the union set of PASCAL VOC 2007 *trainval+test* and PASCAL VOC 2012 *trainval* is used as the training set, and the PASCAL VOC 2012 *test* is used as the test set following [39, 8]. From Table 8, we can also see ARC-FCN still outperforms Faster-RCNN and R-FCN across all object categories.

B. More Qualitative Examples

In addition to the article, we show more qualitative examples in this section.

Fig. 7 shows some qualitative results of ARC-FCN (red) and R-FCN (blue). For the convenience of comparison, the groundtruth annotations are also showed by the green bounding boxes. From these results, we can see our model can localize objects more accurately than R-FCN.

Fig. 8 shows some qualitative results of our model on the PASCAL VOC 2007 testset. We set a relatively high threshold of 0.45 for good visualization, and different object classes are shown with different colors. From these results, we can see our model can fairly well localizing various objects with diverse shapes and imaging conditions.

Fig. 9 shows some typical failure examples of our model. The failure cases are mainly due to occlusion, rough or partial detections and inter-class confusions (e.g., cat and dog, sofa and chair). From the last figure in Fig. 9 and the detection analysis in Fig. 6, we can also see another failure case is the confusion of objects and various backgrounds.

References

[1] O. Aghazadeh, H. Azizpour, J. Sullivan, and S. Carlsson. Mixture component identification and learning for visual recognition. In A. Fitzgibbon, S. Lazebnik, P. Perona, Y. Sato, and C. Schmid, editors, *ECCV*, 2012. 3

[2] H. Azizpour and I. Laptev. Object detection using strongly-supervised deformable part models. In *ECCV*, 2012. 3

[3] S. Bell, C. L. Zitnick, K. Bala, and R. B. Girshick. Inside-outside net: Detecting objects in context with skip pooling and recurrent neural networks. In *CVPR*, 2016. 6

[4] L.-C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille. Semantic image segmentation with deep convolutional nets and fully connected crfs. In *ICLR*, 2015. 6

[5] X. Chen, K. Kundu, Z. Zhang, H. Ma, S. Fidler, and R. Urtasun. Monocular 3d object detection for autonomous driving. In *CVPR*, 2016. 3

[6] X. Chen and A. Yuille. Articulated pose estimation by a graphical model with image dependent pairwise relations. In *NIPS*, 2014. 3

[7] M.-M. Cheng, Z. Zhang, W.-Y. Lin, and P. H. S. Torr. BING: Binarized normed gradients for objectness estimation at 300fps. In *IEEE CVPR*, 2014. 2

[8] J. Dai, Y. Li, K. He, and J. Sun. R-FCN: Object detection via region-based fully convolutional networks. *arXiv preprint arXiv:1605.06409*, 2016. 1, 2, 3, 5, 6, 7, 9, 10, 12

[9] Q. Dai and D. Hoiem. Learning to localize detected objects. In *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on*, 2012. 4

[10] C. Desai, D. Ramanan, and C. C. Fowlkes. Discriminative models for multi-class object layout. *International Journal of Computer Vision*, 2011. 3

[11] D. Erhan, C. Szegedy, A. Toshev, and D. Anguelov. Scalable object detection using deep neural networks. In *CVPR*, 2014. 3

[12] M. Everingham, L. Van Gool, C. Williams, J. Winn, and A. Zisserman. The pascal visual object classes (voc) challenge. *IJCV*, 2010. 3, 6, 7, 9

[13] P. Felzenszwalb, R. Girshick, D. McAllester, and D. Ramanan. Object detection with discriminatively trained part-based models. *PAMI*, 2010. 1, 3

[14] S. Fidler, R. Mottaghi, A. Yuille, and R. Urtasun. Bottom-up segmentation for top-down detection. In *CVPR*, 2013. 4

[15] S. Gidaris and N. Komodakis. Locnet: Improving localization accuracy for object detection. In *CVPR*, 2016. 4, 7

[16] R. Girshick. Fast r-cnn. In *International Conference on Computer Vision (ICCV)*, 2015. 1, 2, 3, 4, 6, 7

method	areo	bike	bird	boat	bottle	bus	car	cat	chair	cow	table	dog	horse	mbike	person	plant	sheep	sofa	train	tv	mAP
Faster-RCNN-Res101	86.5	81.6	77.2	58.0	51.0	78.6	76.6	93.2	48.6	80.4	59.0	92.1	85.3	84.8	80.7	48.1	77.3	66.5	84.7	65.6	73.8
R-FCN-Res101-ReIm†	86.9	84.0	78.4	63.6	59.6	80.5	80.5	91.8	58.6	81.6	58.8	91.1	85.0	85.7	83.7	57.8	81.0	66.3	85.1	72.9	76.7
ARC-FCN-Res101‡	**88.7**	**84.3**	**81.7**	**66.4**	**63.0**	**81.6**	**82.1**	**92.8**	**60.2**	**84.4**	**60.6**	**92.3**	**87.8**	**86.4**	**85.3**	**60.2**	**82.0**	**66.6**	**86.9**	**74.4**	**78.4**

Table 8. Detailed results with IoU $\geq$ 0.5 of Faster-RCNN, R-FCN and ARC-FCN on PASCAL VOC 2012 test set. † : http://host.robots.ox.ac.uk:8080/anonymous/C3H0GM.html. ‡ : http://host.robots.ox.ac.uk:8080/anonymous/WB5KF0.html

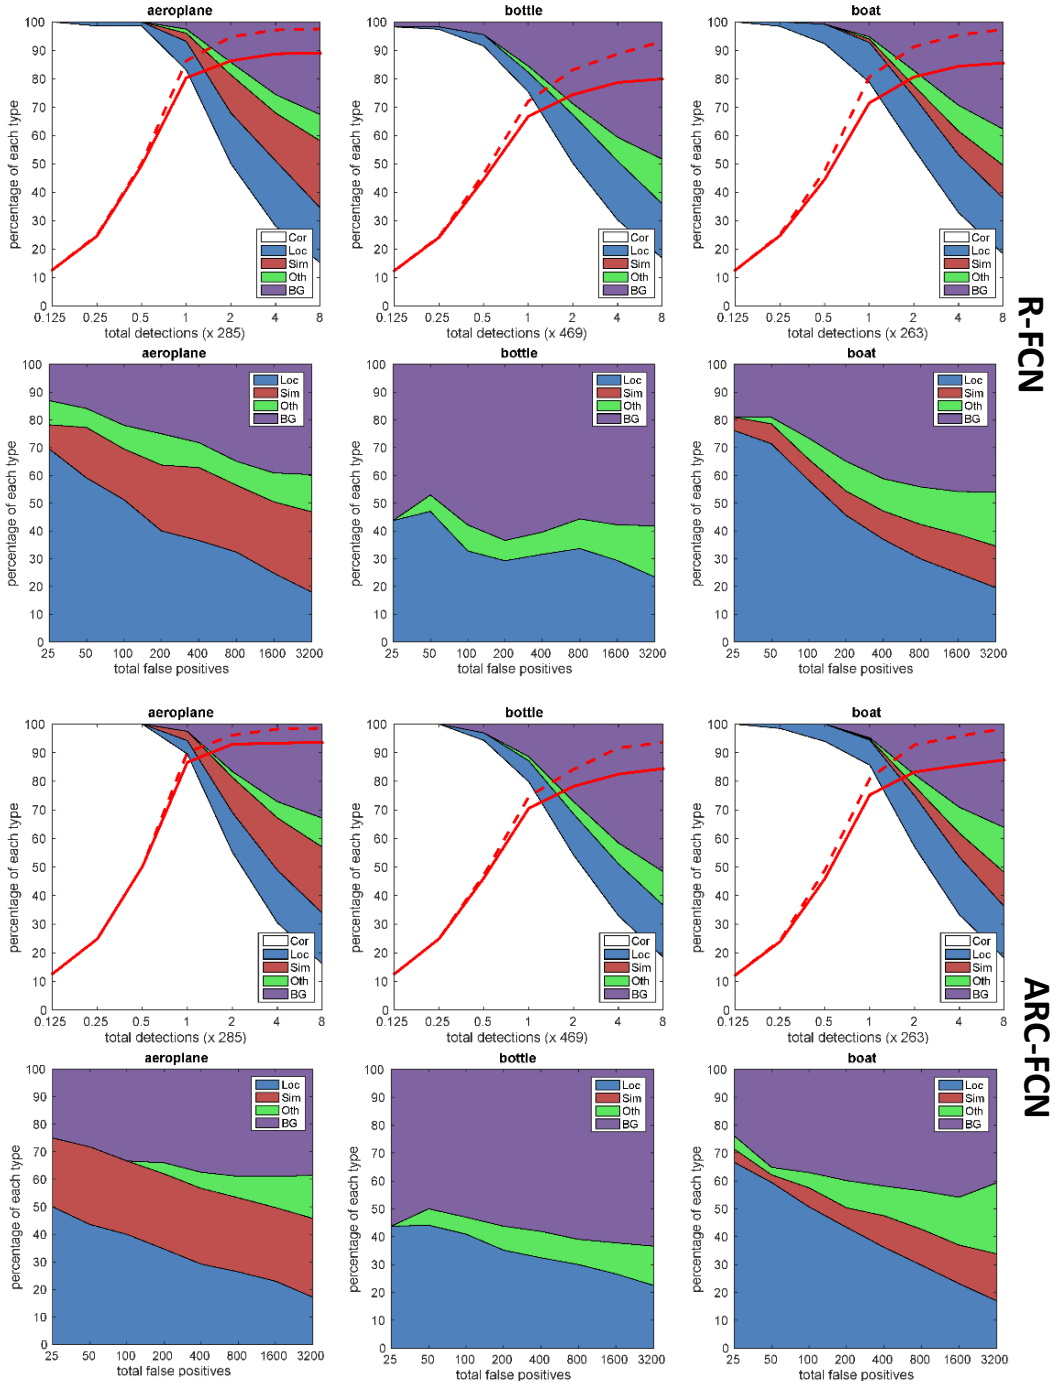

Figure 6. False positive/detection trends with rank. Detection analysis of R-FCN are shown on the first 2 rows. Detection analysis of ARC-FCN are shown on the last 2 rows. The First and Third Rows: type of detection as number of detections increases; The Second and Fourth Rows: stacked area plot showing fraction of FP of each type as the total number of FP increase. line plots show recall as function of the number of objects (dashed=weak localization, solid=strong localization).

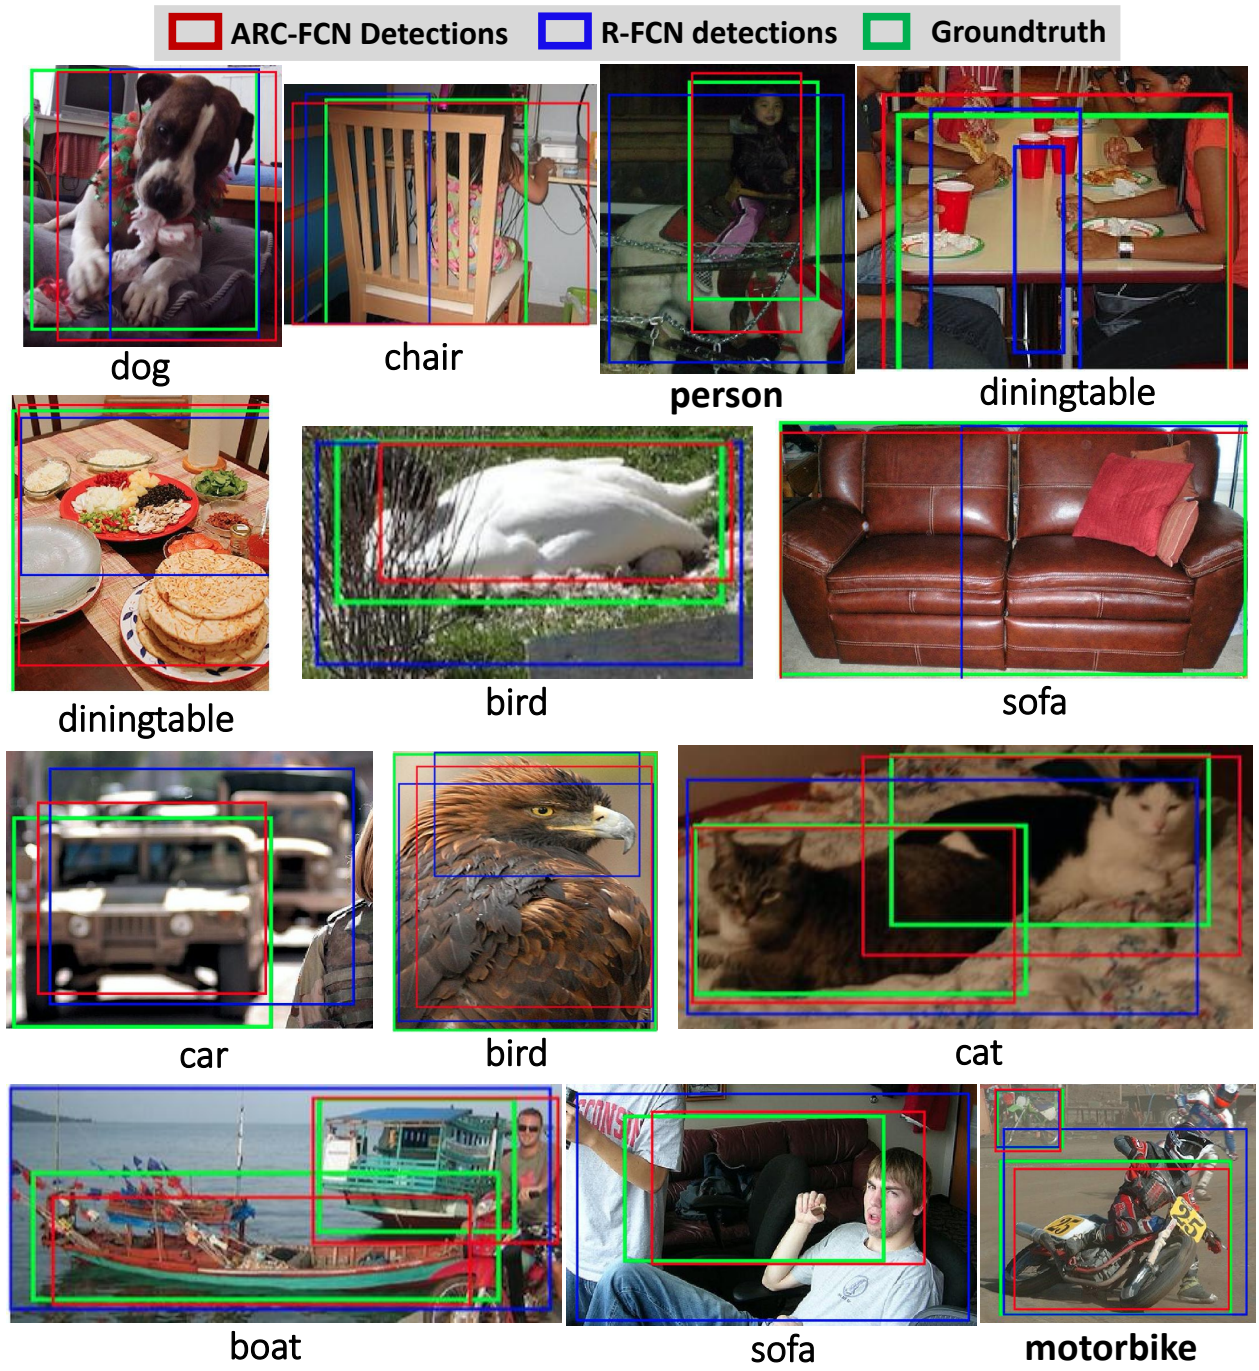

Figure 7. Sample detections of R-FCN-Res101 [8] (blue) and ARC-FCN-Res101 (red). For comparison, we also show the groundtruth bounding box (green). The score threshold is set to 0.6 for good visualization. Best viewed in color and zoom in.

[17] R. Girshick, J. Donahue, T. Darrell, and J. Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *CVPR*, 2014. 1, 2, 3, 7

[18] R. Girshick, P. Felzenszwalb, and D. McAllester. Object detection with grammar models. In *NIPS*, 2011. 3

[19] R. Girshick, F. Iandola, T. Darrell, and J. Malik. Deformable part models are convolutional neural networks. In *CVPR*, 2015. 3

[20] J. X. Guang Chen, Yuanyuan Ding and T. X. Han. Detection evolution with multi-order contextual co-occurrence. In *CVPR*, 2013. 3

[21] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. *arXiv preprint arXiv:1512.03385*, 2015. 1, 3, 4, 6, 7

Figure 8. Qualitative Results of ARC-FCN-Res101 on Pascal VOC 2007 test set. For good displaying, different object classes are shown with different colors. The score threshold is set to 0.45 for good visualization. Best viewed in color and zoom in.

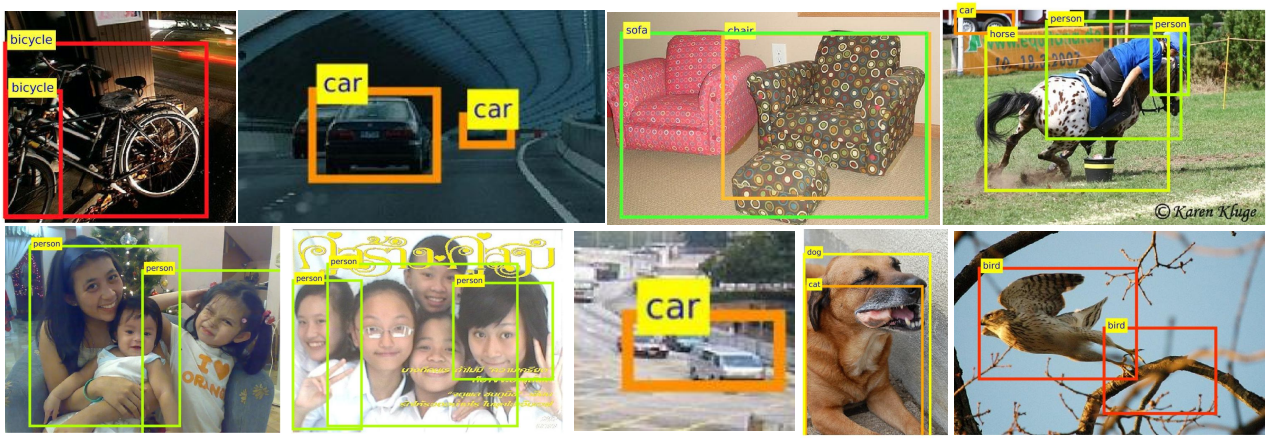

Figure 9. Failure examples of ARC-FCN-Res101 on Pascal VOC 2007 test set. Different object classes are shown with different colors. The score threshold is set to 0.45 for good visualization. The failure cases are mainly due to occlusion, rough or partial detections, and inter-class or foreground/background confusions. Best viewed in color and zoom in.

[22] D. Hoiem, Y. Chodpathumwan, and Q. Dai. Diagnosing error in object detectors. In *ECCV*, 2012. 10

[23] D. Hoiem, A. A. Efros, and M. Hebert. Putting Objects in Perspective. In *IEEE CVPR*, 2006. 3

[24] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell. Caffe: Convolutional architecture for fast feature embedding. *arXiv preprint arXiv:1408.5093*, 2014. 6

[25] H. Kaiming, Z. Xiangyu, R. Shaoqing, and J. Sun. Spatial pyramid pooling in deep convolutional networks for visual recognition. In *European Conference on Computer Vision*, 2014. 3

[26] Y. LeCun, B. Boser, J. S. Denker, D. Henderson, R. E. Howard, W. Hubbard, and L. D. Jackel. Backpropagation applied to handwritten zip code recognition. *Neural Computation*, 1:541–551, 1989. 1

[27] B. Li, T. Wu, and S.-C. Zhu. Integrating context and occlusion for car detection by hierarchical and-or model. In *ECCV*, 2014. 1, 3

[28] H. Li, J. Brandt, Z. Lin, X. Shen, and G. Hua. A multi-level contextual model for person recognition in photo albums. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016. 4

[29] Y. Li, C. Huang, C. C. Loy, and X. Tang. Human attribute recognition by deep hierarchical contexts. In *European Conference on Computer Vision*, 2016. 4

[30] T. Lin, M. Maire, S. J. Belongie, L. D. Bourdev, R. B. Girshick, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick. Microsoft COCO: common objects in context. *CoRR*, abs/1405.0312, 2014. 3, 6

[31] W. Liu, D. Anguelov, D. Erhan, C. Szegedy, S. Reed, C.-Y. Fu, and A. C. Berg. SSD: Single shot multibox detector. *arXiv preprint arXiv:1512.02325*, 2015. 1, 3

[32] R. J. Lopez-Sastre, T. T., and S. Savarese. Deformable part models revisited: A performance evaluation for object category pose estimation. In *ICCV-WS CORP*, 2011. 3

[33] S. Mallat. *A Wavelet Tour of Signal Processing, Third Edition: The Sparse Way*. Academic Press, 3rd edition, 2008. 6

[34] E. Ohn-Bar and M. Trivedi. Learning to detect vehicles by clustering appearance patterns. *TITS*, 2015. 3

[35] W. Ouyang and X. Wang. A discriminative deep model for pedestrian detection with occlusion handling. In *CVPR*, 2012. 3

[36] B. Pepik, M. Stark, P. Gehler, and B. Schiele. Teaching 3d geometry to deformable part models. In *CVPR*, 2012. 3

[37] B. Pepik, M. Stark, P. Gehler, and B. Schiele. Occlusion patterns for object class detection. In *CVPR*, 2013. 3

[38] J. Redmon, S. K. Divvala, R. B. Girshick, and A. Farhadi. You only look once: Unified, real-time object detection. In *CVPR*, 2016. 1, 3

[39] S. Ren, K. He, R. Girshick, and J. Sun. Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. In *NIPS*, 2015. 1, 2, 3, 4, 5, 6, 7, 9, 10

[40] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, 115(3):211–252, 2015. 3

[41] S. Schulter, C. Leistner, P. Wohlhart, P. M. Roth, and H. Bischof. Accurate object detection with joint classification-regression random forests. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2014. 4

[42] P. Sermanet, D. Eigen, X. Zhang, M. Mathieu, R. Fergus, and Y. LeCun. Overfeat: Integrated recognition, localization and detection using convolutional networks. *CoRR*, abs/1312.6229, 2013. 3

[43] A. Shrivastava and A. Gupta. Contextual priming and feedback for faster r-cnn. In *Computer Vision – ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part I*, 2016. 4

[44] A. Shrivastava, A. Gupta, and R. Girshick. Training region-based object detectors with online hard example mining. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 6

[45] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov. Dropout: A simple way to prevent neural networks from overfitting. *J. Mach. Learn. Res.*, 15(1), Jan. 2014. 8

[46] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. E. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich. Going deeper with convolutions. In *CVPR*, 2015. 1

[47] C. Szegedy, A. Toshev, and D. Erhan. Deep neural networks for object detection. In C. J. C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 26*, pages 2553–2561, 2013. 3

[48] A. Torralba. Contextual priming for object detection. *IJCV*, 53:2003, 2003. 3

[49] J. R. R. Uijlings, K. E. A. van de Sande, T. Gevers, and A. W. M. Smeulders. Selective search for object recognition. *International Journal of Computer Vision*, 104(2):154–171, 2013. 2

[50] T. Vu, A. Osokin, and I. Laptev. Context-aware CNNs for person head detection. In *International Conference on Computer Vision (ICCV)*, 2015. 4

[51] L. Wan, D. Eigen, and R. Fergus. End-to-end integration of a convolution network, deformable parts model and non-maximum suppression. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2015. 3

[52] Y. Xiang, W. Choi, Y. Lin, and S. Savarese. Data-driven 3d voxel patterns for object category recognition. In *CVPR*, 2015. 3

[53] Y. Yang and D. Ramanan. Articulated pose estimation with flexible mixtures-of-parts. In *CVPR*, 2011. 3

[54] Y. Zhang, K. Sohn, R. Villegas, G. Pan, and H. Lee. Improving object detection with deep convolutional networks via bayesian optimization and structured prediction. In *CVPR*, 2015. 3, 4, 7

[55] S.-C. Zhu and D. Mumford. A stochastic grammar of images. *Found. Trends. Comput. Graph. Vis.*, 2006. 3

[56] Y. Zhu, R. Urtasun, R. Salakhutdinov, and S. Fidler. segdeepm: Exploiting segmentation and context in deep neural networks for object detection. In *CVPR*, 2015. 3

[57] C. L. Zitnick and P. Dollár. Edge boxes: Locating object proposals from edges. In *ECCV*, 2014. 2